Los Alamos National Laboratory LA-UR-00-1572
Unlimited Release

Electric Industry Simulation System (ELISIMS)

LDRD Project

Report of Progress

March 2000

(This page is blank.)

Electric Industry Simulation System (ELISIMS)

Abstract

The Los Alamos National Laboratory is committed to the development of a Comprehensive, Detailed Simulation of the Electric Power Industry covering all of North America. We include variable resolution adequate to include fine details as required for certain applications and explore the dynamics of the restructured (competitive) marketplace interacting with those traditionally considered in engineering analysis alone. The project name is ELISIMS, an acronym for Electric Industry Simulation System. About two years' effort are reported here: an initial applications study conducted in late FY 1998, a concept and prototype effort in FY 1999, and the beginnings of an intense collaboration with the California Independent System Operator under a Cooperative Research and Development Agreement (CRADA) signed in December 1999. Other collaborations are welcome.

Contributors

L. Jonathan Dowell
Martin Drozda
Dale B. Henderson
Verne W. Loose
Madhav V. Marathe
Douglas J. Roberts

March 2000

Research supported by the US Department of Energy under contract W-7405-ENG. Contact the project leader, Dale B. Henderson, Los Alamos National Laboratory, Los Alamos, NM 87545 or dbh@lanl.gov.

(This page is blank.)

1 Introduction and Motivation

The Los Alamos National Laboratory (LANL) has begun developing the Electricity Industry Simulation System (ELISIMS). The effort enjoys two separately arising imperatives, either of which alone is adequate motivation for a major new undertaking. (1) The recent Presidential Decision Directive 63 (PDD 63) initiates action in response to the report of the President's Commission on Critical Infrastructure Protection, entitled "Critical Foundations: Protecting America's Infrastructures. And, (2) the regulatory infrastructure of the electric industry is already evolving rapidly toward competitive markets in place of the traditional regulated monopolies. Secretaries Peña[1] and Richardson[2] have each voiced strong Department (DOE) and Administration support for this restructuring which reputedly could save consumers $20 billion annually.

At present the electric power industry in the US and worldwide is undergoing restructuring largely in the form of deregulation of generation and re-regulation of the transmission and distribution functions with general movement toward more open and competitive markets. These changes are motivated by shifts in the political and regulatory environments in which it has become reasonable to consider inducing competition. Large state and region-wide electric price differences provide additional motivation. The electric power system has developed over decades under the regulated monopoly paradigm yielding the system we have today, with its stability and security. The system is evolving in response to new market imperatives that will result in a system that may be quite different from that to which we have grown accustomed. Importantly, market-driven initiatives toward free entry and exit of market participants, withholding of capacity in response to current and future price signals, bilateral contracts and distributed generation could lead to a system evidencing a less robust transmission interconnection than we have today.

The Department of Energy (DOE) has clear responsibility under the PDD 63 to work to reduce vulnerabilities within our electrical infrastructure. Doing so requires first knowing exactly what these vulnerabilities are, including interdependencies that lead to subtle vulnerabilities, how equipment failures cascade, and finally how effective possible remedies might be in making the infrastructure more robust. The fact that ownership, management, and control of the infrastructure is rapidly changing because of the current and continuing restructuring toward competition only make a challenging problem more difficult.

[1] "Administration's Plan Will Bring Competition to Electricity, Savings to Consumers; $20 Billion a Year in Savings for Consumers", USDOE News Release R-98-035, March 25, 1998.
[2] "Richardson Releases Administration's Comprehensive Electricity Competition Plan", USDOE News Release R-99-077, April 15, 1999.

The introduction of market incentives via restructuring will cause the emergence of a variety of considerations new to the management of the electric system. In its present form, restructuring has focused primarily on the generation function allowing (indeed, requiring) formerly regulated utilities to divest generation facilities as a step towards introduction of price competition at this stage of the industry. In the newly restructured industry, no one is required to generate to meet demand. Firms may now withhold generation as part of a strategy to maximize profit. The emergence of market participant strategies dramatically increases the complexity of system management.

Recently (4 January 2000) the DOE released the Interim Report of the US Department of Energy's Power Outage Study Team (POST report)[3], in which the lead paragraph reemphasizes these considerations:

> The electric power industry is in the midst of evolutionary change. The reliability events during the summer of 1999 (i.e., outages in New York City, Long Island, New Jersey, the Delmarva [Delaware-Maryland-Virginia] Peninsula, the South-Central States, and Chicago and nonoutage power disturbances in New England and the Mid-Atlantic area) demonstrate that the necessary operating practices, regulatory policies, and technological tools for dealing with the changes are not yet in place to assure an acceptable level of reliability. In a restructured environment, generation technologies and prices are a matter of private choice, yet the reliability of the delivery system benefits everyone. The operation of the electric system is more difficult to coordinate in a competitive environment, where a much larger number of parties are participating.

The intended applications of ELISIMS (section 1.3) include issues highlighted by the POST report.

The huge number of interacting parts in the electric power system makes it extremely difficult to deduce the effect of changes in the system without resorting to computational simulation. The size and complexity of the underlying physical system now coupled with the complexity of markets and strategies requires careful attention to simulation theory and to the underlying mathematics. Vast processing resources will be required to encompass the expanded system. The Accelerated Strategic Computing Initiative (ASCI) of the DOE has provided the LANL with unique computing capability. Pioneering achievements in developing finite event simulations for the Department of Defense (DoD), Department of Transportation (DOT) and other sponsors provide us the expertise and experience to make possible the development and fielding of a *comprehensive*, *detailed* simulation of the electric *industry*:

- *Comprehensive* in that we propose including the whole North American continent because that is becoming the scale of tight interconnection.
- *Detailed* in that we propose to include *each* significant element at the level of generators, transmission, varied control elements, and load distribution buses. *Detailed* also in that market participants whose decisions over operating

[3] "Interim Report of the U.S. Department of Energy's Power Outage Study Team: Findings from the Summer of 1999", USDOE, January 2000.

criteria may be driven largely by profit motives will control each of these physical elements.
- *Industry* in that we intend to include the regulatory, financial, and market factors that interact with the engineering elements.

We have already seen a few cascading failures by which small changes can cause large disturbances over broad areas; these demand *comprehensive* and *detailed* simulation for study and prevention. Deregulation opens the possibility of quickly changing cross-country transfers of power and other industry evolution, emphasizing the importance of the *comprehensive industry* simulation, including more than just the engineering elements. We were all reminded of the importance of this feedback – between market dynamics and engineering dynamics – by the eight 1999 case studies covered in the POST report.

1.1 National Needs

The price of electric power varies markedly across the United States. The prospect of reducing prices by establishing retail access has prompted several states to institute competition in place of regulated monopolies. Electrical power is the last of the regulated industries to be restructured (deregulated) within the past couple of decades in the US. Railroads, trucking, airlines, telecommunications, and banking have all been restructured in recent decades. Some European and South American countries together with New Zealand and Australia have deregulated their electrical power industries. Recently the Secretary of Energy announced that the administration would propose federal legislation to encourage and further enable the trend, estimating that the resulting saving to customers would come to $20 billion per year. Movement away from traditional monopoly regulation will be affected by evolution in generation, transmission, and control technologies that are already changing how the system works.

The security of the nation requires and the DOE is responsible for ensuring that the electrical power infrastructure continues to provide an appropriate quality of service. However, the quality of service will be increasingly determined in markets, with customers defining different grades of quality, depending upon the use to which the electrical power is being put. The Department's responsibility under the PDD 63 following the "Critical Infrastructures" report reemphasizes this. Unfortunately, much of the existing reliability stems from conscious over-capacity in the generation and transmission functions; such over-capacity is naturally threatened by market competition, as service providers are motivated to reduce costs. This point too is reemphasized by the recent POST report. Better, comprehensive simulation will allow study of the necessity for and actual costs of building and maintaining a system that delivers the desired quality of service and reliability. A number of important applications, each serving a national need, are listed and explained in a subsequent section.

1.2 LANL Capabilities

The LANL has much experience and many capabilities to apply to this simulation. Several are unique in their own right; the collection is certainly unique as an ensemble. First, LANL is one of three participants in the DOE Accelerated Scientific Computing Initiative (ASCI) Program; each participant is currently developing a one teraflop computing capability as an approach to the goal of a machine capable of 100-teraflops by 2004. Los Alamos is the designated sole site for the single 30-teraflop capability in 2001. Besides the unique hardware, ASCI (part of the Science Based Stockpile Stewardship effort) are developing the computational knowledge and infrastructure to take advantage of the hardware. These LANL capabilities will be reapplied to the simulation of the electric industry and to other major problems of national significance under a Laboratory effort called Delphi.

LANL is a multi-program defense laboratory of the DOE and is concerned primarily with issues of national security. As such it is properly positioned and equipped to handle security applications of the proposed simulation. From its beginnings during World War II LANL has been a pioneer in phenomenological and physical simulations and associated numerical methods required for successful completion of the proposed simulation. LANL also has decades' experience in military scenario (discrete-event) simulations in which the interactions between people and machines are crucial to understanding the dynamics. Our experience in such "people-in-the-loop" simulation will prove very valuable, especially with the important extensions to market applications including the participants' competitive strategies and tactics. LANL has experience in air quality simulations, realistically modeling complex terrain in detail, including the Mexico City air shed with its manifold pollution sources and mountainous topography. This experience will become relevant as dispatch calculations include environmental considerations along with economics, stability, and security. LANL even has a small in-house generating station and substations, providing some convenient local expertise in utility operations.

As a federally-owned laboratory with no commercial interests and with much experience in handling classified materials, LANL is well-suited to serve as an honest broker and custodian of confidential data which may be required for certain applications.

Unique LANL capabilities include the Transportation Analysis Simulation System (TRANSIMS) project sponsored by the US Department of Transportation (DOT). This project has developed and applied microsimulations at the individual traveler level, integrating multiple independent (but interacting!) instances together with the roadways and signal network to effectively simulate traffic patterns. TRANSIMS has been successfully demonstrated in a study of a 25 square mile area within the Dallas-Fort Worth metroplex including 200,000 vehicles over a five-hour span. TRANSIMS and other LANL developed simulations of complex adaptive infrastructures provide experience and methods to be applied to the comprehensive, detailed electric power industry simulation.

The TRANSIMS project has also spawned an expertise in developing synthetic populations at the household level that form the basis for activities that then can be used to derive a transportation demand function. This maintains all important statistical variation and properties to support simulations, but the data are synthetic in that they are derived from actual census data but do not use actual census data that, in any case, are not available due to privacy and confidentiality considerations. An obvious extension is to use our synthetic population methods to drive the demand for residential electric power.

Over the past decade LANL has developed a capability to evaluate the U.S. electric-power grid. We have broad expertise in using graph-theory and power-engineering models to analyze the electric-power grid. In addition, we have conducted significant data collection and data mining efforts to establish an extensive database describing the U.S. electric-power infrastructure. Sources of data include Federal Energy Regulatory Commission, Energy Information Administration, North American Electric Reliability Council, state Public Service Commissions, commercial databases, and electric power utilities. We have conducted analyses that identify chokepoints, key components, and the consequences of component failures.

Under Electric Power Research Institute (EPRI) sponsorship LANL has experience in studying the computational complexity of executing various possible clearance algorithms in competitive retail electricity markets. One important conclusion of this work is that some possible policies for bilateral contract satisfaction pose basic problems that are computationally intractable even when restricted to very simple power networks[4]. (They are NP-hard problems.) As a result it is unlikely that one can implement some proposed policies exactly and efficiently for actual networks. This motivates the use of powerful computational resources (as are being developed in the DOE's ASCI program) to solve these problems approximately.

1.3 Varied Applications

We foresee varied important applications for our comprehensive detailed simulation of the electric industry. Because there are so many applications we have structured this section into seven categories:

1. Policies and Regulation
2. Infrastructure Security
3. Technology Evaluation
4. Data Evaluation
5. Dispute Resolutions and Litigation

[4] D. Cook, V. Faber, M. Marathe, A. Srinivasan and Y. J. Sussmann, "Combinatorial Problems in Production and Transmission of Electric Power: Theory and Experimental Results." Technical Report LA-UR-99-0048, 1999, Los Alamos National Laboratory, Los Alamos, NM, 87545.
D. Cook, G. Hicks, V. Faber, M. Marathe, A. Srinivasan, Y. J. Sussmann and H. Thomquist, Experimental Analysis of Contract Satisfaction Problems Arising In The Deregulated Power Industry, Preliminary Version, July 1999.

6. Software Verification and Validation
7. User Facility

Within a given category, the questions to be studied can have a number of disciplinary flavors. A technically motivated policy might be aimed at improving the margin of transmission network stability. A market-oriented policy might establish competition at the household level by allowing small bilateral contracts. A financial policy might impose a tax to "buy out" stranded costs. Each of these might serve its intended purpose fully and yet cause serious unintended consequences as seen from another discipline. Our ultimate intention is to serve each of these communities of interest with integrated, consistent tools. Most importantly, this consistency facilitates the study of the couplings between them. While unintended consequences cannot be eliminated, our comprehensive test bed will allow its users to anticipate as many as possible.

1. Policies and Regulation. The *Comprehensive, Detailed* Simulation of the Electric Power *Industry* will be a crucial tool with which to test the practicality and efficacy of new policies and regulations being considered within or promoted to the DOE and its Federal Energy Regulatory Commission (FERC) or to various state, provincial or local authorities. Possible new policies and regulations warrant careful analyses for a variety of reasons. For example, how should costs, which have been bundled under the traditional vertical monopolies, be separated in the new competitive markets? Or, to what extent is electric power (energy) a commodity and to what extent is it a service? Competition encourages a commodity-trading model while the digital economy stresses reliability and power quality, which are services. Unbundled costs might be assigned differently under the commodity and service models, yet both models apply today. In deriving costs for (unbundled) ancillary services, the industry and the government should employ detailed models to discover where and how the costs are incurred and to better assess the effects of failures.

New policies can have unintended consequences. Sometimes these are technological. For instance: to what extent might the "Million Solar Rooftops" of small urban photovoltaic generators compromise network security or power quality as shifting cloud patterns bring them erratically out of service? Could these fluctuations cascade into larger effects distant from the rooftops in question? Might a sudden onset of clouds in Sacramento cascade into a brownout in New Jersey?

Market policies necessarily affect people who are very good at responding with efficient, but unforeseen dynamics. Anticipating unforeseen consequences requires that the simulation include market dynamics – sometimes resolved to the single household level – together with power network engineering. Much of the unscheduled demand is at the household level, emphasizing the need for synthetic population methods.

2. Infrastructure Security. The assessment of the security of the electric infrastructure commands greater attention, even without the uncertainties introduced through restructuring, and ever more so because of it. The recent Presidential Decision Directive 63 (22 May 1998) and the corresponding report of the President's Commission on

Critical Infrastructure Protection (entitled Critical Foundations, October 1997) emphasize the threats from both hostile and accidental causes. The DOE and the nation have basic security issues to address, many of which would best be undertaken with our proposed simulation. A secure environment like that available at LANL would be required to protect and control sensitive and/or proprietary information.

States, municipalities, the Federal Emergency Management Administration, the military and other agencies might also utilize the simulation capabilities to study contingencies, natural catastrophe and accident scenarios, and possible terrorist vulnerabilities within their jurisdictions. Because of the interdependencies of the continental network, these could portend even broader consequences than those initially anticipated.

3. Network Control Technology Evaluation. Possible new technologies pose important new issues for analysis. Because a new gizmo could be developed and deployed does not always mean that it should be. Beyond questions of whether it will really work as intended when deployed on a large scale, there are attendant questions of the effects on market economics, system reliability, system vulnerability to cascading failure, and issues of infrastructure protection (possible unintended new vulnerabilities due to the new equipment).

For example, network control will naturally become distributed as autonomous control and switching elements are fielded. These devices will be more and more computer controlled, requiring that the deployed software itself be fully tested in the context of the distributed environment. Such testing will demand *literal* incorporation of deployed (and proposed to be deployed) software within the comprehensive, detailed electric industry simulation. Executing such a simulation complete with literal distributed operational software imposes a great computation challenge, appropriate to the DOE ASCI environment.

4. Data Evaluation. Questions of the utility and value of data are going to escalate. The DOE and its Energy Information Administration require ever more detailed data from operating companies. With growing competition the reluctance to surrender such data, even under safeguards of confidentiality, will escalate. It would be valuable to know which data are important, for which purposes, and to what extent. Again, our comprehensive, detailed simulation of the electric industry would provide invaluable guidance in proposed rule making and in the drafting of necessary legislation.

Less contentious perhaps, but equally important are evaluations of possible new engineering data. New technologies provide the opportunity for new sensors and transmission of more and more real-time data. Operators need to understand the value of further data, under variants of distributed, centralized, and hybrid control regimes. This needs to be explored in detail before expensive new sensors are deployed. Estimating the cost of developing new sensor technology may be possible in isolation, but estimating the operating system derived worth from having the data available requires the comprehensive, detailed simulation.

5. Dispute Resolutions and Litigation. Each of the above four topic areas is likely to generate contention and litigation. A comprehensive, detailed simulation of the electric industry and the expertise required to use it could function as an informed means to resolve controversy before litigation ensues. If we move into an open market at the retail level in which it is impossible to satisfy all bilateral contracts, new and better-defined notions of economic fairness will be required. The system-wide implications of these differing notions of fairness will then lead to policy issues, dispute, and possibly litigation. We have previously explored these issues on a smaller scale and would continue them as part of the fifth application area.

6. Software Validations and Verification. Several vendors already offer software packages for unit commitment, economic dispatch, power flow, etc. These are purchased by and useful to the operating utilities that can conveniently operate the simulations on their own affordable computing platforms. In the future markets for these products will expand as at least the major industrial customers of electricity feel the need to be better-informed market participants. Such lesser simulations will remain useful and necessary. But, just how good are they as the overall network moves to operation with reduced margins? Under some basis to be determined by the sponsor, our comprehensive detailed simulation could be used to test and calibrate these more aggregated simulations to allow their confident utilization where appropriate, and to avoid it where dangerous.

7. User Facility. LANL could perhaps serve as a kind of user facility to operating utilities, power producers, system operators, power exchanges, and to major consumers. Again the problems posed would couple technology with market strategy. They would cover varied time scales: one for price response, another for a load-shedding strategy, and yet another for transient stability. Confidentiality of commercial studies would have to be maintained, as an "honest broker." With no vested interest and with experience in handling classified information, LANL could do this very well. Rules for availability and possible compensation to the government would have to be developed by government sponsors.

Each of these seven application areas alone could be argued to merit the development of the comprehensive detailed simulation of the electric industry. Together they define a national scale capability worthy of Department cognizance and of major government investment.

1.4 Diverse Users

The varied applications grouped into the seven areas above are each important to a broad set of users who will belong to government, to electric industry providers, to investors, to members of academia, to customers and to the public. The levels of detail (devices to transmissions lines to markets – or mixed) and the scope (local to continental) will vary with the application. So too will the required urgency of response: some results will have lasting value, others will demand various degrees of urgency because a response too late may be of no value.

Policy-oriented questions will be important to federal agency and department personnel who make rules, to industry personnel who must abide by regulations, and to academic and public interest people who wish to assess likely consequences. Can retail wheeling be made to work at the household level? Would it result in cascading failures? Would the open competitive market somehow learn to avoid failures because failures benefit no one? Questions such as these couple the levels of simulation (e.g., markets and ac power flow) together and become computationally challenging. But these sorts of computation results need not be delivered in real time, rather only with adequate timeliness to avert human user frustration.

Questions assessing vulnerabilities are important to operators, to government response planners, to the military, and to other authorities. Contingency analyses are important to operators. But the questions are related. One agency may need to know what ensues should a terrorist dynamite a line pole, another may wonder about a truck running into a pole by accident; the ensuing failures or the system robustness against failure may be the same. Both groups may employ “red team” users to probe vulnerabilities. Because vulnerabilities may indeed be discovered with the simulation, controls on its use may be required or wise. The security afforded by the Los Alamos site may become important.

Contingency analysis may sometimes arise when the contingency is at hand. Because the response times required of system operators are short, it would be very difficult to set up, execute, and forward detailed analyses of results from an unforeseen eventuality. Therefore, real-time applications on the scale of system transients are not part of our intended set of applications. Nevertheless, the design of the simulation system, its databases, and its other inputs could facilitate or further limit urgent applications. It is important in the design phases to ensure that the maximum practicable responsiveness be maintained.

Optimization applications, of interest to system operators, require real-time responses if the network is to be optimized to current conditions. But these real-time responses may afford more time than operating contingencies; we should be able to accommodate them sooner.

Questions of technology evaluation will be of interest to operating utilities and technology developers considering investment. For example, would a new device under consideration really prove effective when distributed over the network? Possible inventions are, of course, intellectual property of value and investment strategies are proprietary. The noncommercial nature of the laboratory may be important to users studying possible new technologies.

The question of whether items of data are legitimately useful to whom, and for which purposes will enter discussions between the owners of such data (e.g., system operators) and regulators or public interest personages arguing for publication. Since useless effort can be expended releasing or protecting data of no consequence, it may be important to each to study system scenarios with varied data sets. Operating companies, customers, regulators, and public interest groups may be very interested; but these results will not be time urgent.

The processes of restructuring the industry are sure to lead to contention and litigation as discussed above. The achievement of a competitive market with physics-based system constraints is likely to lead to failed contracts, leading to arbitration and litigation. The comprehensive, detailed simulation of the electric power industry is likely to be a useful tool for litigants, arbitrators, and the courts. Again the applications would not demand real-time responsiveness, but the impartial, non-commercial, laboratory with a tradition of security could be important.

Software validation and verification would entail comparing (for given scenarios) results from our comprehensive detailed simulation with those from other (generally commercial) software packages. Both the vendors and users of such packages, which could run on an operator's computer, would be interested in such results. Both have an abiding interest in determining under which conditions convenient commercial packages should and should not be used. Again the unbiased Los Alamos environment is important in offering these results.

1.5 POST's Specific Action Items

Our seven broad application areas in section 1.3 resulted from our FY 1998 study. They formed the motivating rationale for launching the ELISIMS project. Following up on its technical findings in its Interim Report[3], the Power Outage Study Team (POST) concluded with a dozen "recommendations to enhance reliability" in its March 2000 Final Report[5]. We note that our seven applications map well into the Final Report's "12 recommendations, each of which includes specific action items for federal consideration"[6]. The relationships are mapped in Table 1 (on following page) with symbols ·· for a primary application and · for a secondary application.

2 Technical Approach

We propose to develop an integrated set of simulation capabilities – a "family" – with which to support a wide range of applications by a range of users. This ambition is warranted, is practicable, and is challenging. It is challenging because simulation projects with all-encompassing goals often succeed in none. But it is equally true that limited initial goals often prove limiting forever, that software is usually not expansible beyond its original intent because of myriad consistent design decisions which are difficult to discover and prohibitively expensive to remove. Fortunately we have experience in avoiding this dilemma.

[5] "Final Report of the U.S. Department of Energy's Power Outage Study Team", USDOE, March 2000.

[6] "Energy Department Issues Recommendations to Help Prevent Power Outages", USDOE News Release R-00-068, March 13, 2000

Our proposed first step was to completely specify and document the manifold applications sketched out above. In doing so we have attempted to brief and listen to as many potential users as we could identify. Help and support from the DOE, from various agencies concerned with national infrastructure security, from EPRI, and from some operating utility companies was invaluable. This definition of the ultimate application space is documented to define both the breadth and depth of proposed applications. The breadth part will be narrative text, an expansion of that already begun with our sections here “Varied Applications” (above) and “Simulation Scope” (below).

	Recommendations from the POST’s Final Report / ELISIMS’ Applications areas	Policies and Regulation	Infrastructure Security	Technology Evaluation	Data Evaluation	Dispute Resolution	Software Validation	User Facility
1	Promoting market-based approaches to ensure reliable electric services	• •	• •	• •			•	
2	Enabling customer participation in competitive electricity markets	• •		• •		•	• •	• •
3	Removing barriers to distributed energy resources	• •	• •	• •				
4	Supporting mandatory reliability standards for bulk-power systems	• •	• •	• •				
5	Supporting reporting and sharing of information on "best practices"					•	• •	• •
6	Enhancing emergency preparedness activities for low-probability, high-consequence events on bulk-power systems		• •	• •	• •			
7	Demonstrating federal leadership through promotion of best reliability practices at federal utilities							•
8	Conducting public-interest reliability-related research & development consistent with the needs of a restructuring electric industry	• •	• •	• •	•	•		
9	Facilitating and empowering regional solutions to the siting of generation and transmission facilities	• •		• •	• •	• •		
10	Promoting public awareness of electric reliability issues							
11	Monitoring and assessing vulnerabilities to electric power system reliability		• •	• •	• •			• •
12	Encouraging energy efficiency as a means for enhancing reliability	•	•	•			• •	• •

Table 1. Mapping of 12 POST Recommendations into ELISIMS’ Application Areas

In order to avoid the trap of a narrow initial design, we will develop the simulation architecture for the eventual ensemble, defining the complete (ultimate) architecture as best we can. In practice this preliminary design will mature piece-by-piece as specific parts are implemented and are refined and extended. We conduct periodic design reviews to ensure that the overall design, its class hierarchy, and the proposed class libraries all appear practicable and to cover the enumerated applications. This aspect of the review is to defend the goals of “breadth”.

In order to avoid the converse trap of beginning too much and never completing anything, we have defined a subset capability for initial prototype implementation. This was completed in FY1999. With guidance from sponsors and users we defined an illustrative application and the corresponding software classes and undertook implementation. Details of the prototype are discussed in section 5.2. In December 1999 the Laboratory signed a Cooperative Research and Development Agreement (CRADA) with the California Independent System Operator (Cal ISO). Our FY2000 development implements those portions of our design focusing on the CRADA tasks, as noted in section 5.3. Subsequent development will continue in similar targeted phases or "blocks". While power flow is a well-developed sub-discipline, we must make extensions to address bigger problems on multi-processor parallel computing architectures; these algorithmic concerns are discussed in section 4.1.

The point of this dual approach is simple. If we do not clearly envision the comprehensive capability in the beginning, we will not achieve it. If we do not begin implementation with a clearly defined "first bite", a particular application, we will not achieve anything in particular.

2.1 Simulation Scope

As already mentioned, we intend simulation at the scale of (literally "of") the North American continent. The oceans on the east and west coasts of the continent provide a natural limit to the problem. At present one could probably get by with simulation on the scale of a major interconnection area of which there are four on the continent. In designing software there is little merit in compromise only to a factor of four. One can always execute sub-problems as is expedient, but the design ought to be toward the entire continent.

We intend resolution down to the level at which each element can be adequately described with a few parameters for all electrical and mechanical (i.e., inertial) components. These include generators, transformers, transmission lines, switches, breakers, FACTS (Flexible AC Transmission Systems) gear, phase angle regulating transformers, and reactance. How many such components exist? The published enumerations vary, but we anticipate:

- 6,000 to 17,000 generators,
- 50,000 to 140,000 transmission lines,
- 40,000 to 100,000 substations, and
- 130,000,000 customers.

These dimensions pose a truly staggering combinatorial problem. While even with the 100-teraflop ASCI machine we are not likely to handle the indicated problem anytime soon, we will be able to undertake significant subsets – far larger ones than possible now. Doing so will enable important new applications that become crucial as the nation moves toward a deregulated, competitive electric industry.

Does this mean that we will jump to modeling 130 million toasters all at once? Clearly not. But implementing a simulation family with variable resolution means that we will be able to include as many detailed devices interacting together and with as broad a context as is required by the problem at hand. This multi-resolution approach is essential to our eventual complete ELISIMS system; it is addressed in section 2.2.

For many simulations relevant to competition, the millions of small-scale customers can be aggregated, with our synthetic population techniques, into loads from the substations. This is justified in that individual, small customers are only poorly characterized at present and are connected radially to the system (that is, without complicating loop flows) from load buses at substations. But considering only the first three lines of the above enumeration still defines a big application!

Questions of proposed regulations and the effects on (say) market power could probably be addressed with such "power flow" calculations further simplified through linearization into the so-called "dc equivalent" problem. (Power-flow computational technology is reviewed in section 4.1.) Questions of network stability and "voltage collapse" are becoming more important as transmission systems are more heavily loaded with long distance "wheeling" of power. These complex power calculations require reactive as well as real power flow, making this simulation broader in scope as well as much larger in scale than current capabilities. Finally, stability calculations require the dynamics of rotating machinery, further expanding the number of variables to be included.

Some market proposals involve freely associated bilateral contracts across long distances. The recent study undertaken at Los Alamos for EPRI is an example[7]. Precise market calculations or studies of the effects of market decisions on network operation and robustness require customer models. These would utilize our experience with synthetic, but entirely representative populations, but would probably be implemented with only appropriate subsections of the continental network. Certain statistical properties of these simulations would then be incorporated into the aggregated loads used in the major simulation.

To better serve the variety of applications listed above, we anticipate a flexible "family" of simulations capable of being instantiated in a variety of modes with consistency from one to the next. We foresee "dc equivalent" power flow, complex power flow, and stability calculations with and without individual customer effects across the entire continent or consistently limited to smaller areas. Members of this "family" of simulations will be linked into a hierarchy of levels as is discussed in the following section.

Some applications enumerated above are more-or-less traditional engineering problems only on a grander scale. We believe that all of the electrical engineering characterizations are well understood and available, except possibly for the only lightly applied emerging technologies, such as FACTS. The scale of the calculation and linking the levels into a

[7] See footnote 4 for references.

hierarchy impose severe questions of computational science and numerical methods. New methods will have to be invented, developed, or reapplied from other domains. Our success here will be dependent upon the fortunate co-location of simulation science being developed in the far larger ASCI effort.

Other applications impose more novel interdisciplinary questions. How difficult is it, computationally, to implement a possible economic or market policy? What is a quantifiable measure of fairness? How computationally complex is the implementation of a given “fair” market policy? What extensions might be required to usefully employ game theory in the satisfaction of competitive, but constrained, markets? Fortunately we have an interdisciplinary team of engineers, physicists, mathematicians, computational experts, and economists, between whom we can address these questions consistently.

2.2 Simulation Hierarchy

Our ultimate objective is a set of simulations linking differing levels of resolution. The simulation at a given level would be represented as a network of links between nodes. This obviously reflects the physical structure of the system, but also allows us to employ a formal hierarchy. Doing so allows us to employ the application of important graph-theoretic (mathematical) techniques, including simplicial complexes and graph grammars. Entities at a given level will be linked to corresponding ones at the more-resolved and more-aggregated levels to facilitate transitions as appropriate. Examples of levels would include the short-time stability level, the complex power flow level, the “dc-equivalent” power flow level, the market level, and the financial level. These levels differ in their degrees of resolution or aggregation, in the spatial extent of an application problem, and in their time scales.

The networks underlying the several levels each capture the interactions between elements or agents represented at that level. As a result the networks need not be identically connected; they are not isomorphic. And the elements represented will differ. Our complete design will have to specify these elements or agents, their dependencies or interactions, and the rules of interactions. The interactions between engineering entities derive from physics; those from market or financial entities derive from economics; those between are less well developed.

Existing network simulations, including those at Los Alamos, illustrate the concept at a given level, power flow for example. Developing consistent representations across levels is more difficult, but we have done so in other contexts. Mathematical and computational issues arise. These include questions of representation: which processes and details are actually needed at a given level? We will consider equivalence between simulation systems and computational efficiency: given two equivalent simulations, which is computationally more efficient? The resolution of these issues requires mathematical theory including computability, dynamical systems, automata theory, and logic.

3 Collaboration

Very early in our effort, in April 1998, we shared our concepts with key staff of EPRI in Palo Alto, enjoying much encouragement, many suggestions, and valuable critique. EPRI contributions color some of the application ideas discussed in section 1 of this report. Collaboration with EPRI staff affords us immediate industry perspective, valuable information, and continuing invaluable criticism.

Beginning in late 1998 we undertook collaboration with the Energy Systems Research Center (ESRC) of the University of Texas at Arlington (UT-A). We have obtained the complete source for a non-linear or AC power flow code developed by ESRC/UT-A faculty and students. The product of many years' development in legacy Fortran, this code offers us an excellent starting point for the implementation of nonlinear power flow in a modern (parallel multi-CPU) environment. It is logical for us to undertake this development in continued collaboration with them.

Our closest collaboration began in early 1999 with the coincidental, but fortunate visit of the General Counsel of the California Independent System Operator (Cal-ISO) to Los Alamos. (Under terms of the industry restructuring begun in March 1998, the Cal-ISO is a not-for-profit corporation charged with system operation, reliability and security and with market surveillance.) Several subsequent technical interchanges lead to the conclusion of a formal DOE-approved Cooperative Research and Development Agreement between Cal-ISO and the Laboratory being signed in December 1999. Being right in the midst of vexing complex issues arising through the coupling of technology with market dynamics, the Cal ISO is ideally suited to advise us on near term priority applications and to supply domain expertise to enable our joint successful solutions. Our current (FY2000) effort is focused on the CRADA applications.

Two of the necessary steps toward the CRADA were formal designations of "unclassified" for our software and the acquisition of an unrestricted export license from the Department of Commerce, although we are not planning on exporting anything.) Both of these formalities facilitate other possible collaborations, which we continue to welcome.

4 Architecture of the Simulation

A flow schematic for ELISIMS is shown in Figure 1 (on the following page). The key feature of our architecture at this level is the clear and separate encapsulation of the various market, engineering and simulation-engine entities: the boxes in the chart. The market entities are generally along the left side, the power flow engine is the central engineering function, and the simulation control is at the top. The control dynamics and principal data flows are indicated by the various arrows; the controlling feed-back loops are evident.

We employ a master and slave design, with parallel distributed slaves as determined by the computational requirements. Available parallelism (not shown in the figure) includes executing several disparate modules in parallel, multiple cases of a given module in parallel, and/or distribution of a single computational module across multiple computers. The concept is, of course, evolving; the figure is our early 2000 version.

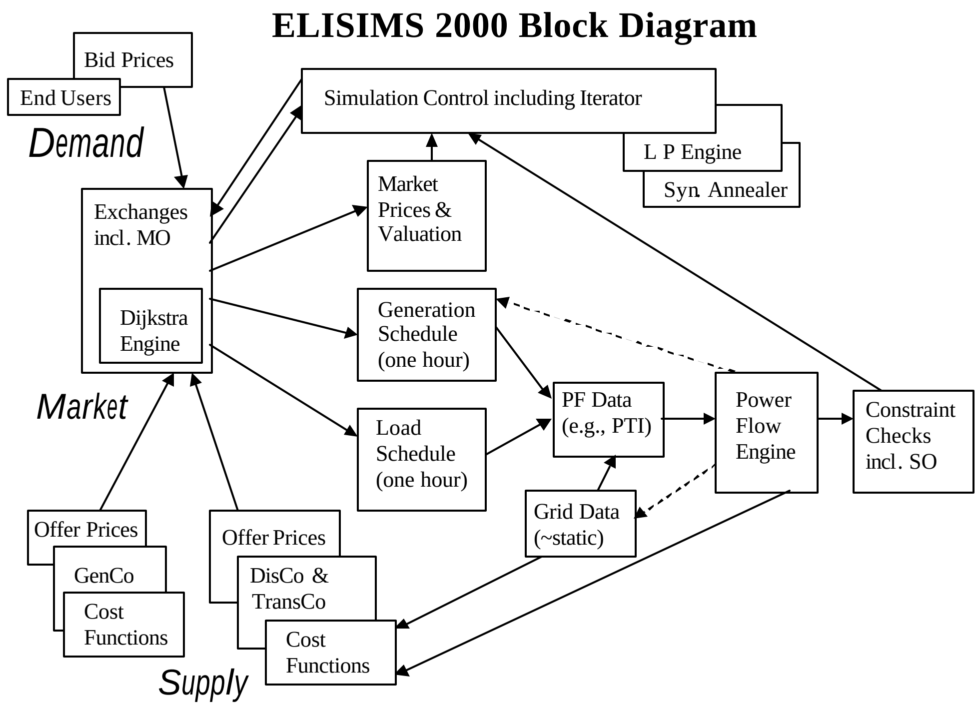

Fig. 1. ELISIMS Architecture, circa early 2000.

Under the concept of a family of simulations, a given instantiation need not (and in general will not), utilize all of the modules shown. The reader will notice this in the detailed description of our 1999 prototype, which follows in section 5.2. Other important general remarks include the observation that a given box in the architecture may have different implementations with different degrees of complexity as appropriate for different applications. The power flow equations, for example may be nonlinear and complex (the so-called full "AC" model), they may be linearized and decoupled into real and imaginary (reactive) power or they may be linearized to the dominant terms (the so-called "DC" model). (There are also highway traffic analogues, which we do not employ.)

4.1 Power Flow Engine

Although the science is well understood[8], and several methods are now described in standard texts[9], the power flow module will continue to be computationally demanding. Therefore we review power flow a little more fully. For most purposes, it is possible and convenient to reduce the common three-phase circuits to ("positive sequence") single wire equivalents. These are usually nondimensionalized on a "per unit" basis in which static transformers disappear and voltage magnitudes are normally near unity. At each node, called a bus in power circuits, the sum of all of the current and power are zero. Generally the flow though the lines connecting to the bus is considered apart and we speak of "injected power" from generators; loads are represented as negative injections. (Reactive power supplied from a generator or a shunt connected to a bus can be of either sign.)

For a line segment between buses i and j, the (complex) current I_{ij} is simply the voltage difference between the buses divided by the impedance, or (equivalently) multiplied by the admittance:

$$I_{ij} = (E_i - E_j) \;/\; Z_{ij}; \qquad \text{with} \quad Z_{ij} = R_{ij} + j\,X_{ij} \qquad \text{or}$$

$$I_{ij} = Y_{ij} \bullet (E_i - E_j); \qquad \text{with} \quad \begin{aligned} Y_{ij} &= G_{ij} + j\,B_{ij}; \\ G_{ij} &= R_{ij} / (R_{ij}^2 + X_{ij}^2), \\ B_{ij} &= -X_{ij} / (R_{ij}^2 + X_{ij}^2) \end{aligned}$$

To construct the equations for (complex) power flow leaving a bus i down the line to bus j, one multiplies the complex conjugate of those for current flow by the voltage at that bus:

$$P_{ij} + j\,Q_{ij} = E_i \bullet Y_{ij}^{\,*} \bullet (E_i - E_j)^*;$$

Summing for the total Power, $P_i + j\,Q_i$, at bus i yields the convenient standard matrix form:

$$P_i + j\,Q_i = E_i \Big[\sum_{\substack{j=1 \\ j \neq i}}^{N} Y_{ij}^{\,*} \Big] E_i^* + E_i \Big[\sum_{\substack{j=1 \\ j \neq i}}^{N} (-Y_{ij}^{\,*}) \; E_j^{\,*} \Big].$$

$$= E_i \quad M_{ij} \quad E_j^*$$

---

[8] N. Sato and W. F. Tinney, "Techniques for Exploiting the Sparsity of the Network Admittance Matrix, IEEE PAS 82, (1963) 944-950.

[9] A.J. Wood and B.F. Wollenberg. *Power Generation, Operation and Control.* John Wiley and Sons, 1996. J. J. Grainger and W. D. Stevenson, Jr., *Power System Analysis* McGraw-Hill, New York, 1994

In practice the P and Qs are (mostly) specified and the matrix problem is solved for the Es – the vector of voltage magnitude and phases. The matrix is clearly very sparse because the off-diagonal elements may be seen to represent lines connecting the nodes. In the WSCC, for example, the number of lines in a typical representation is only about 1.3 times the number of nodes.

The matrix M_{ij} is obviously symmetric in occupation in that a connection from *i* to *j* is also a connection from *j* to *i*. The connections through simple lines (implied above) are truly symmetric, but others – notably phase-shifting transformers introduce asymmetries: $Y_{ij} \neq Y_{ji}$. The simple derivation given above also implies that the real and imaginary parts of M are diagonally dominant: a consequence of $M_{ii} = -\sum M_{ij}$ and of the Xs being mostly of one sign (inductive). However, this is not quite true because there are lines compensated with series capacitor banks (which contribute negative line reactances) and – more commonly – because shunts connected to buses contribute to the diagonal elements without appearing in the off-diagonal terms. Even so, the matrix M_{ij} *almost* has nice properties. This "almost nice" nature of the matrix is exploited in some methods of solution.

4.1.1 Methods of Solution

The Gauss-Seidel method consists of repeated application of the non-linear equation for complex power, $P_i + j\,Q_i$, – node by node -- using an admixture of newly obtained and old values, iterating across the whole problem until the (complex) values for E_i converge. Data-dependent variable devices, such as on-load tap-changing transformers and controlled shunts, are allowed to change with the data, dynamically modifying the impedances (the matrix elements) during the course of iteration.

The Newton-Raphson method[10] utilizes an iterated first-order linearization using the Jacobian matrix, the terms of which are $\partial P/\partial\theta$, $\partial P/\partial V$, $\partial Q/\partial\theta$, $\partial Q/\partial V$, where V and θ are the magnitude and phase of the voltages E and the subscripts have been omitted. Again the iteration continues until the Es appear to converge and again the dynamic devices are allowed to modify the matrix elements.

Of the two methods, Newton-Raphson converges more quickly, but Gauss-Seidel is more tolerant of poorly chosen (guessed) initial data, especially for the phases. Many authors advocate starting with one method and switching to the other. For this reason we have implemented the Gauss-Seidel method into the UT-A legacy code which employs Newton-Raphson. We have demonstrated that both methods converge to the same answers, even with on-load tap-changing transformers included. But we are yet to demonstrate any real advantage from the added flexibility.

[10] W. R Tinney and C. E. Hart, "Power Flow Solution by Newton's Method," *IEEE Trans. on Power Apparatus and Systems,* 86 (1967) 1449-1460.

The reactances, X, typically dominate the resistances, R. The voltage magnitudes are typically near unity, so the voltage differences are dominated by the differences in phases. Making these approximations yields the linearized or (so-called) "DC" approximation:

$$P_{ij} = (\theta_i - \theta_j) \; / \; X_{ij},$$

which is obviously much simpler. (One needs to notice that in this – the "DC" approximation – an alternating current effect – reactance – is dominant.) We have found it helpful to first solve a linear approximation to establish initial values for the phases and then switch to the non-linear problem.

In between the full "AC" treatment and the "DC" approximation, decoupled methods are sometimes employed. These are based on the observation that real power is mostly dependent upon phases and the reactive power is more strongly affected by voltage magnitudes. That is, one uses the dominances:

$$| \partial P/\partial\theta | \;>>\; | \partial P/\partial V | \quad \text{and} \quad | \partial Q/\partial V | \;>>\; | \partial Q/\partial\theta | \qquad (\text{ with } |V| \approx 1)$$

to discard two quarters ($\partial P/\partial V$ and $\partial Q/\partial\theta$) of the Jacobian and to decouple the problems for P and Q. We have not employed these methods to date.

Large problems appear to be much more difficult to solve than smaller ones. This, coupled with our need to employ multiple CPUs to a single large problem, have prompted us to investigate various methods of partitioning the matrix problem. The literature contains many variants of "diakoptics", meaning the tearing of a problem into pieces[11,12,13]. There are two motivations for and two separations of the problem: one to enable the computation to be shared among several processors, and one to take advantage of having an "almost nice" sparse matrix.

Our initial trials have been encouraging[14], both with methods employing "tearing" along physical boundaries of the network and by separating the "almost" nice matrix into a sum of a "nice" matrix and lesser-rank one with the offending terms, that is methods based on

---

[11] G. Kron, "Diakoptics - piecewise solution of large scale problems," The Electrical Journal, in 20 partss from 158, 1673-1677 (1957) to 162, 131-141 (1979).

[12] J. R. Bunch and D. J. Rose, "Partitioning, tearing, and modification of sparse linear systems," J. Math. Anal. and Appl. 48, 574-593 (1974).

[13] P. W. Aitchison, "Diakoptics as a general approach in engineering," J. Eng. Math. 21, 47-58 (1987).

[14] L. J. Dowell, "How Diakoptics Works: An Example Power-Flow Problem," LANL Technical Report LA-UR-99-6847 (Dec. 1999).

the inverse of sums of matrices. In no case, however, is it practical to actually invert the large sparse matrix.

4.2 Energy Market Models

The model must represent the behavior of market entities as they interact in the market for electrical energy. Specifically, generators, generation suppliers, transmitters, transmission suppliers, and distributors of electrical energy will enter the market as suppliers of the service, asking prices for supply. Other entities, who may enter bilateral contracts, or who have ownership in supply purchased from one of the asset-owning market entities may also enter the market and behave much like suppliers, even though they have no ownership of generation, transmission, or distribution assets. End-users enter the market to acquire price information and make commitments for supply of electrical energy. Arbitrageurs may represent end-users or aggregations of end users. The market operator must integrate asking prices, new contracts struck at current market prices and limit prices to determine a new market price at the time any of these variables change. The market modules must be capable of interfacing with a power flow module to transfer information about detailed schedules for various components of the system.

4.2.1 Principles of Market Design

It is generally accepted that the functioning of markets should evince the following principles – (i) permit market participants to promote utility maximization, (ii) promote economic efficiency, (iii) insofar as possible, delegate the decision-making process to the buyers and sellers, (iv) ensure that the information structure (information availability and transmission) should be such that (a) it does not yield undue economic advantage to any one party (market power) and (b) it is complete and yet minimal to the extent possible, and (c) the information structure should be such that it can be calculated and transmitted efficiently to all the market players. Traditional operating principles adhered to by the power industry are

- Meet predicted time varying demand at minimum operating cost
- Compensate for transmission losses (real and reactive)
- Meet various operating constraints such as voltages at buses, thermal ratings of transmission lines, etc.
- Provide flexible generation in real time to balance deviation from anticipated demands
- Maintain secure and reliable operations by maintaining standby network resources

The advent of the deregulated power industry implies that the operating principles in the new environment may be significantly different. In particular, the system operator will be given the responsibility of maintaining a secure and reliable operation; the market economics will govern prices and generation of power. As a result, the market players have no direct responsibility or ability to adhere to the traditional operating principles listed above.

4.2.2 Assumptions

Our design assumptions include

- Any entity may enter the market to bid or offer generation, transmission flow, or end-use consumption
- Service quality (assumed to be firm) and specific location defined by consumption node (for end-users) or generation node (for generators)
- Electric energy must be delivered from generation nodes to consumption nodes using the existing transmission infrastructure. The network infrastructure has all the basic engineering parameters defined (e.g., capacity, maximum voltage rating, etc.)
- Bilateral contracts for electricity and/or its delivery are permitted

Contracts for generation, transmission flow and distribution flow are consummated by the party's acceptance of the current market price.

5 Simulation Entities

The software entities indicated in Figure 1 correspond to real-world entities. In keeping with the principles of object oriented software design, we carry the entity modularity as deeply as is practicable. This results in typical simulation entities having representations both in the market model and in the physical (power flow) model; we keep these two representations separate. For instance, a generator (machine) has well defined physical characteristics included in full power flow codes, such as that from UT-A.

In solving for a consistent power flow, a generator is usually represented as a "PV" bus, meaning that the real power injected into the net and the bus voltage maintained by the generator are specified. The reactive power, Q, injected from the generator (which may have either sign) will be whatever is required for a consistent solution; similarly the phase of the generator will be that which makes the power flow algorithms converge. However, the power flow code will also monitor the P and Q together to ensure that limiting characteristics of the physical device are not violated. In some cases the real power output is adjusted to accommodate network losses, in these cases the power flow package must ensure that the real power found is consistent and feasible for the machine.

As a market entity, a generator (an owner) has a business strategy and has tactics. A generator might simply accept a price set by outside entities and provide whatever power is sought, between zero and the physical limits. Or he might publish a table listing varoius amounts of power for sale at different prices, trying to better match his customer demands with his particular operating characteristics. Or he may want to withhold power from the future market in order to “make a killing” because of shortfalls and corresponding higher prices in the spot market. Or he may favor the stability of locking in long-term bilateral contracts. Or perhaps he finds it profitable to not sell real power, but instead sell ancillary services needed by the system operator to maintain system security. The range of business strategies is wide, varied, complex, and interacts with the physical quantities.

For each simulation entity – generators, transmission, end-users, etc. – we must develop appropriate physical and market representations. The mature ELISIMS will maintain families of each.

5.1 California Protocols

In early 1999, we formed an abstraction of the California system from a variety of published sources, mostly on the internet sites of the California Power Exchange (Cal PX) and the California Independent System Operator (Cal ISO). This abstraction is summarized in Figure 2. The figure shows how various information flows apply to the day ahead-market, the first of five stages in the complicated California protocol.

First, the cross-over between the aggregated demand bid curve and the aggregated supply bid curve results in the Unconstrained Market Clearing Price (UMCP). This is published by the Cal PX which conducts the underlying auction.

Next, balanced schedules (in which total supply match total demand) are submitted by the Schedule Coordinators (SCs). The Cal PX is one SC, other brokers and traders submit balanced schedules for which the prices need not be made public. From the total flow implied by all of the schedules, the Cal ISO calculates the net loads on the various transmission circuits. If there is no congestion (no constraints are violated), the a priori schedules are to be implemented.
If congestion results from the aggregation of the balanced schedules (i.e., if constraints are violated), the Cal ISO then utilizes Schedule Adjustment Bids (both positive and negative) from the same SCs to adjust flows until all of the loads are served and no constraints are violated. Substantial sums may be exchanged as a result of this process as the resulting differences in zonal market clearing prices become transmission prices.

Similar processes are employed for the same day market, using Incremental Bids. Real-time stability and security of the system are achieved by the Cal ISO utilizing Supplemental Bids as necessary. If the Schedule Adjustment, Incremental and Supplemental Bids are insufficient for the Cal ISO to assure operating security, it may impose other dispatch schedules upon the generators as necessary. The Cal ISO also conducts markets in the Ancillary Services – Automatic Generation Control, Spinning Reserve, Ready Reserve, Standby Reserve – as are also required to maintain system security, generally on a real-time basis.

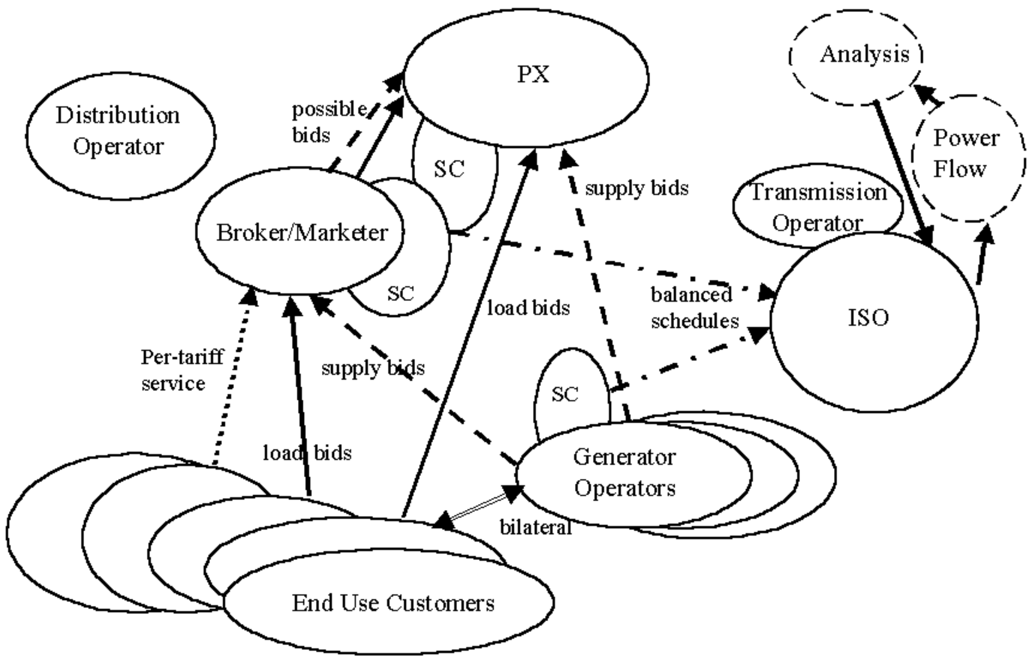

Fig.2 Abstraction of California Protocols for Day-ahead Market

All of this is complex and is still evolving. Therefore the California Protocols were not employed within our initial prototype. The understanding achieved in making the abstraction will, however, benefit us in implementing the (current) California protocols as part of our joint CRADA effort with the Cal ISO.

5.2 Components of the Prototype Simulation

The complete ELISIMS requires much more than power flow, integrated into a package with feedback mechanisms reflecting the emerging real work coupling market dynamics with electrical dynamics. For this larger application we must abstract the models to be simulated and invent methods to drive them. The team debated the relative merits of developing a simulation that would mimic the physical attributes of a specific “real world” (specifically California) electric system together with its market protocols versus developing a hypothetical system with generalized attributes applicable to a typical system. Both possible simulations were considered at some length. Either way nothing would be lost because ELISIMS is required to implement multiple markets coexisting upon, and interacting through, a single engineering infrastructure.

For the prototype, the latter strategy was the one selected. The rationale for this choice is that the purpose of the simulation at this stage of implementation was to demonstrate the feasibility of interfacing a market model with a model of the physical system and to perform cursory investigations of the dynamic interplay with the two systems. Being a functioning prototype, it also supports timing studies that enable us to better estimate the computational resources required to support various applications. Using a hypothetical market – rather than California – on the prototype also obviated any possible questions as to the accuracy or adequacy of the abstraction from reality.

5.2.1 The Power Flow Infrastructure

The electric power infrastructure developed for the prototype simulation is a relatively simple network containing 4 generation nodes, 18 consumption nodes, and 44 links, or transmission lines, connecting the generation and consumption nodes. While small enough to execute on a laptop computer, this turns out to be big enough to demonstrate some interesting results.

Initially the prototype included the UT-A legacy code with its full nonlinear "AC" treatment. Later we implemented the simpler linearized power flow, which is more appropriate to the simplified problem, resulting in much faster performance.

5.2.2 The Prototype Market Model

The market model developed for the prototype is relatively simple in concept and is intended primarily as a placeholder for more comprehensive models that will be incorporated into the ELISIMS project at a later date. It does nevertheless demonstrate strong coupling between market and engineering functions using transmission (congestion) price as feedback signal. It also demonstrates interesting behavior typical of a sequential dynamical system.

5.2.3 The Prototype's Simulation Entities

The simulation entities employed in the prototype reflect real-world entities in both the physical and the economic realms.. Some objects have more than one representation (generators, for example, supply power to the physical model and prices to the economic model). Referring back to the schematic architecture of figure 1, we enumerate the software entities instantiated for our prototype below:

1. **Generators:** A decision-making entity owning or controlling physical assets. Each generation unit is a separate decision element.

A generator's market representation consists primarily of a production schedule for the 24-hour futures period. For each of the 24 hours, the generator posts a schedule of power available for sale that hour, and the price that the generator will charge to supply that quantity of power to the grid.

Obvious possible extensions include generators with strategic logic: let the generators post prices that vary as the market for a given (forward) hour progresses.

2. **End-Users or Consumers:** A decision-making entity that consume electrical energy; enters the market to bid price for supply.

   A consumer's market representation in the prototype model is comprised of a schedule of the consumer's power requirements for each hour of the 24-hour futures market, and a preference to purchase the least expensive power available.

   Obvious possible extensions include purchasing the rights (options) for more power than an end-user expects to need. The extra is then later re-sold at a premium, re-sold at a loss, or defaulted which incurs a monetary penalty.

3. **Transmission Operator (TO):** A decision-making entity that owns or controls physical transmission assets and asks a price to carry electric flow across specific line segments. Each line segment contains a bus at each end.

   The transmission operator posts prices that vary as the market progresses in response to the line loadings as published by the System Operator (SO). The objective is to employ price feedback to preclude congestion. (In the end the SO will void offending sale contracts if necessary; the intent of the market design is to make this unnecessary.)

4. **Transmission Line:** The transmission network, with its constraints, is required for the power-flow calculations.

   (Clearly the generator function (item 1) should be reported here as a market function and an engineering function also; the point of displaying it above as a single entity is to provide a contrast, illustrating the distinction.)

5. **Market Operator (MO):** The Market Operator accepts a (in the prototype, the) set of consumer bids and is responsible for brokering contracts with individual generators on behalf of the end-user clients. The MO has two objectives:

   a) Procure contracts that will satisfy each consumer's power requirements for each hour of the futures period, and

   b) Procure these contracts on a least-cost basis.

   To accomplish these objectives the MO performs the following computational steps for each hour the futures market:

   - Creates a random ordering of the customers. This emulates the consumers entering the market at random times. Rerunning the prototype with different random number sequences results in interestingly different results for the end-users and for the generators.

- Choose the first customer from the random ordering. Compute the least-cost generator for that consumer. Dijkstra's algorithm is used to compute the "contract-path-price" linking the consumer and each available generator. An available generator is one that still has power available for contracting for the current bidding period. The price of delivered power to the consumer is computed as the generator point-of-origin production price for the current hour, plus the sum of all transmission line access fees, plus any transmission congestion fees.
- Choose the generator with the lowest contract-path-price. If the generator has at least as much power available as the consumer's present requirement, request that the SO approve a contract in that amount between the consumer and the generator. If the generator has less power available than the consumer's requirement, request the SO approve the partial amount contract.
- If the SO disapproves the contract, pick the next least expensive generator and attempt to satisfy the consumer's power requirement. Repeat this process until either the consumer's requirements are satisfied, or there are no more feasible generators.
- Select the next random customer.

The bidding process for the futures market is declared closed when the MO has processed all customers. Note that at the conclusion of the market, all customer demand may or may not be satisfied. The prototype allows frustration: early congestion can preclude further sales with both unsatisfied demand and available power for sale.

6. **System Operator (SO):** The SO runs the power flow code to assess transmission congestion in order to approve provisional transmission contracts. The SO also calculates true power flows ("loop flows") on the line segments. The SO is a market entity with three responsibilities:
   - Determine if a contract being requested by the MO would violate any system constraints, such as a transmission line capacity,
   - Compute the new transmission line utilizations after each contract has been approved, and
   - Report the true loads to the TO who then computes a new price after each incremental contract is struck.

   If in this process a line constraint is found to be violated, the proposed contract is limited in amount. The MO must then try the next best contract on offer to meet the balance of the particular end-user's demand.

We conclude this section with some observations about the prototype model:

1. The market uses a nodal pricing rule—price to deliver power into the system at each and every node. Accordingly it is possible, indeed likely, that prices will be different at different nodes. This is due both to different generation prices at

different nodes and also to different line loadings, and hence differing congestion charges on different lines.

2. After the SO approves every contract, it uses the new line utilization for computing the next contract-path-price. Because of this, load is continuously preferentially shifted to lesser-utilized transmission lines because they will have a lower add-on congestion fee. Also, the delivered futures price of power to each consumer is continuously updated.
3. A *least cost path* based method is used to compute the price a consumer will pay for the power received. The price calculations include the production cost, a fixed transmission cost, and a variable transmission charge that depends on the line utilization of the network.
4. Each consumer employs a *least cost* strategy for entering into a bilateral contract. In this strategy the user chooses a generator that can provide the required power at the least cost basis. The consumer demand is assumed inelastic and no re-trades are accommodated.
5. The generator does not have any influence (apart from setting its price) over the way it is chosen to be a party in a bilateral contract. Moreover, as in the case of the consumer, the generators cannot re-trade a contract that is agreed upon and the amount of power they can supply cannot be changed to reflect subsequent market conditions.
6. Market frustration is possible. Certain congestion conditions can preclude further contracts, even with unsatisfied customer demand and a power available for sale. This mechanism could be explored for very interesting strategic behaviors by the market participants.
7. The outcomes vary markedly with the sequence of random choices. This is typical of sequential dynamical processes, of which our prototype marketplace is an example.
8. Besides the possible extensions already mentioned, it would be easy to add a concluding loop in which one-by-one the contracts are removed and restored to the final (settled) market results. The true marginal prices thus discovered could be compared to the sequentially determined contract prices. The degree of agreement would be a measure of market efficiency.

5.2.2 Restatement of the Prototype Model Algorithm

This section restates in sequence the algorithm that implements the market and physical models described in the previous section.

Step 1: Generators post a schedule of availability and price for each hour of the forward market day.

Step 2: Customers post a schedule of their demand for power by hour for each hour of the forward market day.

Step 3: For each customer, one-by one, in the random ordering, the MO:

i. Computes the least cost available generator capable of supplying all or part of that customer's demand. If the least cost generator for a given customer can supply only part of the customer's demand, the MO finds the next least costly generator capable of making up the difference in that customer's demand.

ii. Communicates the provisional order to the System Operator (SO) who then runs the power flow code to determine that no line loading constraints are violated by the provisional order being completed. If no line loading constraints are violated the SO approves the order and notifies the MO.

iii. Computes the customer's charge for the power generated and the delivery of that power to the customer's place of consumption. The delivery charge (posted by the TO) includes a basic line charge and a congestion charge that is related exponentially to the load on that line.

iv. Subtracts the generator availability just allocated to the last customer from the availability remaining at the close of the prior transaction.

Step 4: The TO posts new transmission prices, reflecting the true loads. The price feedback to the market is intended to cause subsequent customers to avoid heavily loaded circuits.

Step 5: The MO steps through each customer in the random ordering in the manner of Step 3 until all of the customer demand has been met or until all of the generation availability has been allocated. Note that it is possible to have either unsatisfied demand or excess generation capacity.

5.3 The California ISO CRADA Market Model

Our FY200 effort is focused on delivering the capabilities outlined in the CRADA with the Cal ISO. The details are to be mutually defined as summarized in the following points. A complete description will be prepared jointly, after the first year of collaboration.

1. The subject area of the project will be the Western Systems Coordinating Council (WSCC) domain, as broken down into various control areas.
2. CRADA Task 1 is the study of congestion management in the WSCC. The first phase of this effort will be concerned with developing methods and software.
3. We will be concerned with the forward market, conducted as 24 separate one-hour auctions. The software design should allow a future extension into more, smaller increments.
4. Generators will be represented in the auctions as supply bids (marginal prices) in the California Power Exchange (CalPX) format. The data will be provided by the Cal-ISO to reflect historical records under various network conditions.
5. Demands will similarly be in the CalPX format and supplied by the Cal-ISO.

6. The resolution (degree of aggregation) of both supply and demand bids will be as specified by the Cal-ISO.
7. The transmission constraints and costs will be specified by the Cal-ISO. Initially these data will be resolved only to the level of inter-ties between the control areas.
8. Using Cal-ISO specification of the CalPX algorithms, LANL will develop a simulation software to clear the markets, producing a Market Clearing Price (MCP) in three variants: (a) taking the entire WSCC as a single market, (b) taking each control area as a separate market, and (c) taking specified groups of control areas as separate markets.
9. Using the power injections and loads from item 8, LANL will develop a power-flow simulation to calculate the loads on the control area interties. Initially this may be done with the "dc" or linear approximation.
10. Congestion from overloads on these interties will be detected and reported. The LANL-developed simulation will than mitigate the congestion using various (alternative) algorithms specified by the Cal-ISO. The resolution will be initially at the control area, therefore the current Cal-ISO constraint on market separation between Schedule Coordinators will be relaxed as appropriate.
11. Using the above data and algorithms, the LANL simulation will report unconstrained and zonal (by control area) market clearing prices (MCPs). Using the transmission prices and the congestion management algorithms, the simulation will report locational prices.
12. The LANL simulation will report the changes in dispatch and imports resulting from each stage of and each variant of the Cal-ISO specified congestion management algorithms. The changes in zonal (congestion area) prices and in locational prices will also be reported.
13. The LANL simulation will report the distributions (by control area) of <u>undispatched</u> generation resulting from the various alternative market and congestion management policies. These data will be used by the Cal-ISO to estimate the adequacy of (potentially) available ancillary services at each control area, under the various market and congestion management policies.
14. Future extensions of this (LANL & Cal-ISO) collaboration will develop more sophisticated treatment of network security as supported by the ancillary services as well as exploring separate explicit ancillary service markets.